OPENFABRICS ALLIANCE

PCI Express* 3.0 Features and Requirements Gathering for beyond

Debendra Das Sharma
Senior Principal Engineer, I/O Technologies

Intel Corporation
April 5th 2011

Legal Disclaimer

- **INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL® PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. INTEL PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.**
- **Intel may make changes to specifications and product descriptions at any time, without notice.**
- **All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.**
- **Intel, processors, chipsets, and desktop boards may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.**
- **Nehalem, Westmere, Romley, and other code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user**
- **Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.**
- **Intel, Intel Inside, and the Intel logo are trademarks of Intel Corporation in the United States and other countries.**
- **Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.**
- ***Other names and brands may be claimed as the property of others.**
- **Copyright © 2011 Intel Corporation.**

PCIe* Roadmap and Trends

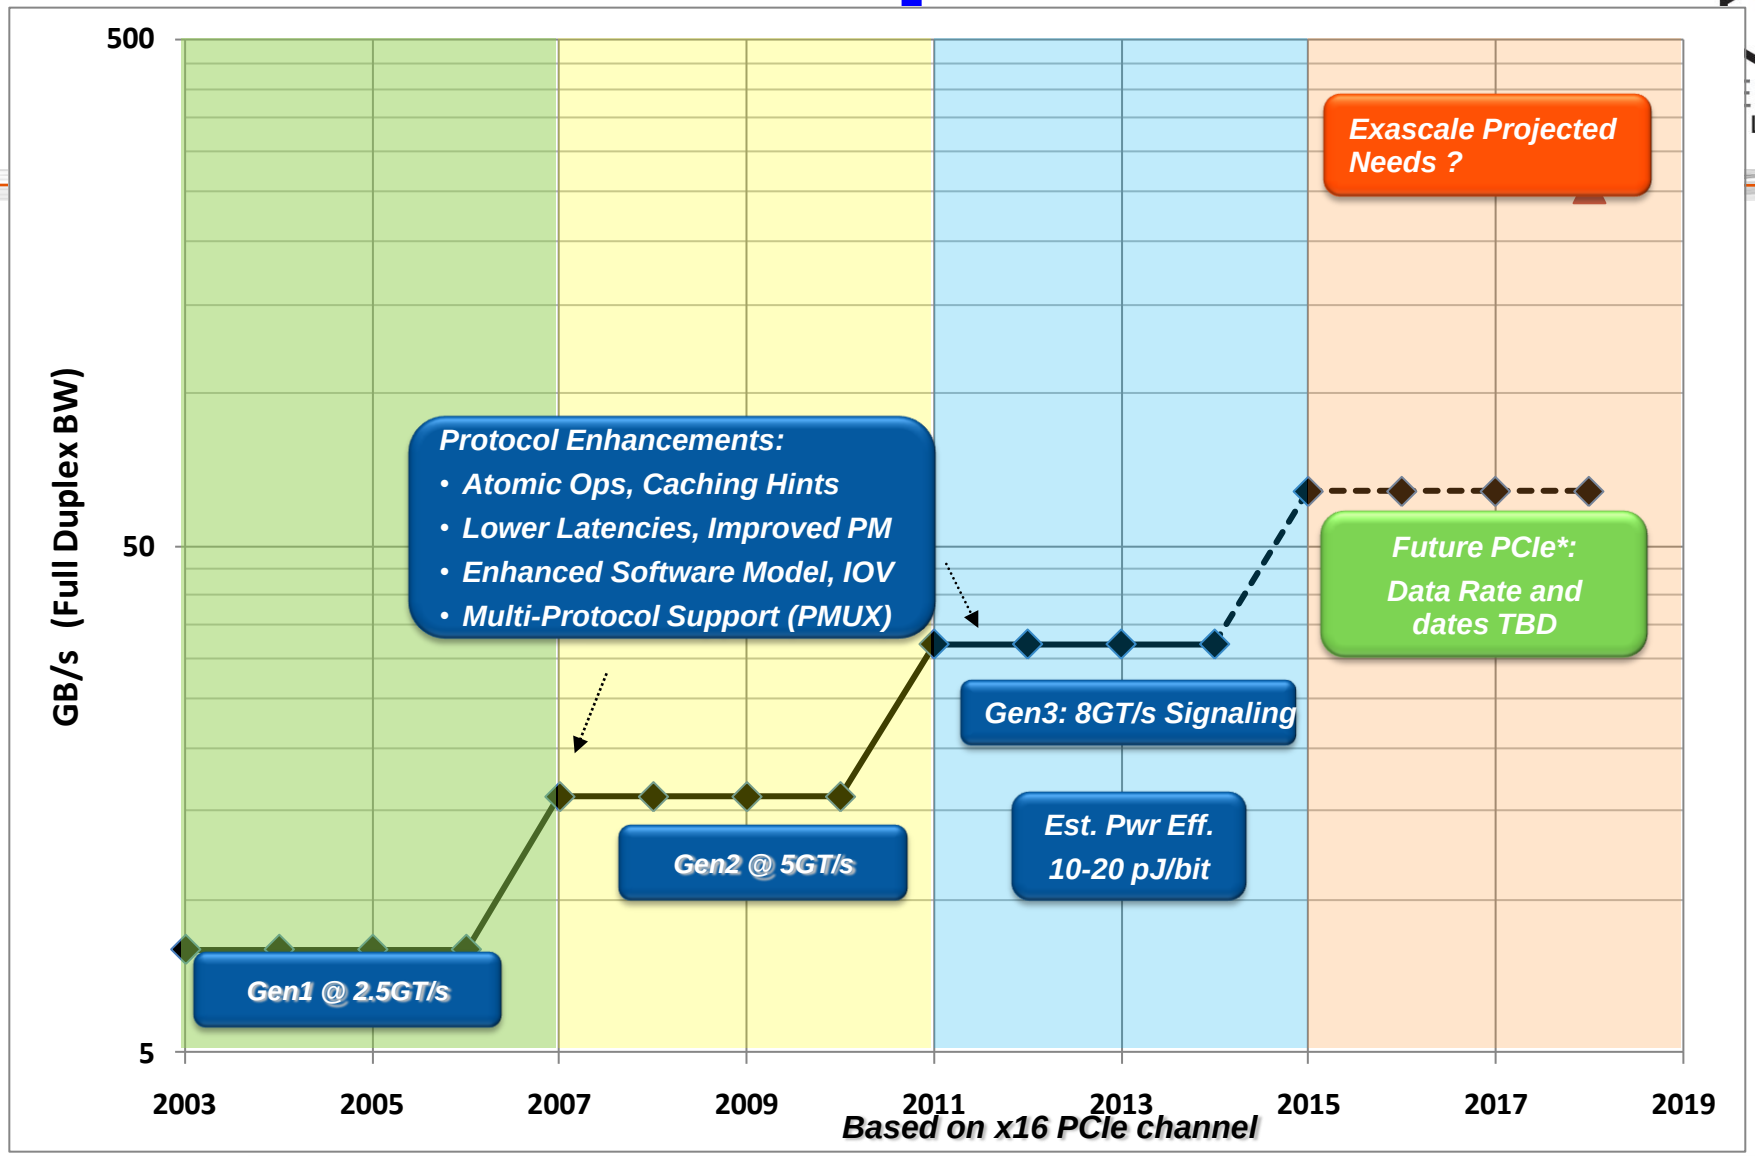

Continuous Improvement: Data Rate as well as Usage Models. Doubling Bandwidth & Improving Capabilities Every 3-4 Years

PCIe* 3.0 Requirements and Data Rate

- **Compatibility with PCIe* 1.x, 2.x**
- **About 2x performance bandwidth over PCIe* 2.0**
- **Similar cost structure (i.e. no significant cost adders)**
- **Preserve existing data clocked & common clock architecture support**
- **Maximum reuse of High Volume Manufacturing (HVM) ingredients**
  - FR4, reference clocks, etc.
- **Strive for similar channel reach in high-volume topologies**
  - Mobile: 8", 1 connector
  - Desktop: 14", 1 connector
  - Server: 20", 2 connectors
- **Evaluated both 8 GT/s and 10 GT/s - 8GT/s met all these requirements**
  - 8 GT/s provides only 60% improvement in B/W
  - A further 25% efficiency improvement towards doubling the b/w comes from replacing the 8b/10b encoding with a new 128b/130b encoding scheme

Backwards Compatibility, HVM capability, Power efficiency, and wide range of topology support were important

Requirements beyond PCIe*3.0?

Considerations:

- Data Rate (Bandwidth): 2X,3X (i.e. 16, 24 GT/s)
- Cost: Modest Cost Increase
- Compatibility: Desirable to be Compliant with PCIe 1/2/3
- Power: ISO Power with PCIe* 3.0
- Channel: 1 Connector 12”; repeaters/cables for longer channels
- Reduce Transaction overhead

Channel Assumptions:

- Mid-loss board material with improved connectors
- Pad-Pad Insertion Loss limited to ~25dB
- Other Possible changes: Back-drill vias/ micro-vias, Circuits: Reduced pad cap, more DFE, reduced voltage swing for low power (backwards compatibility implications)

Lead Applications: HPC, Networking and Storage

Options beyond PCIe*3.0

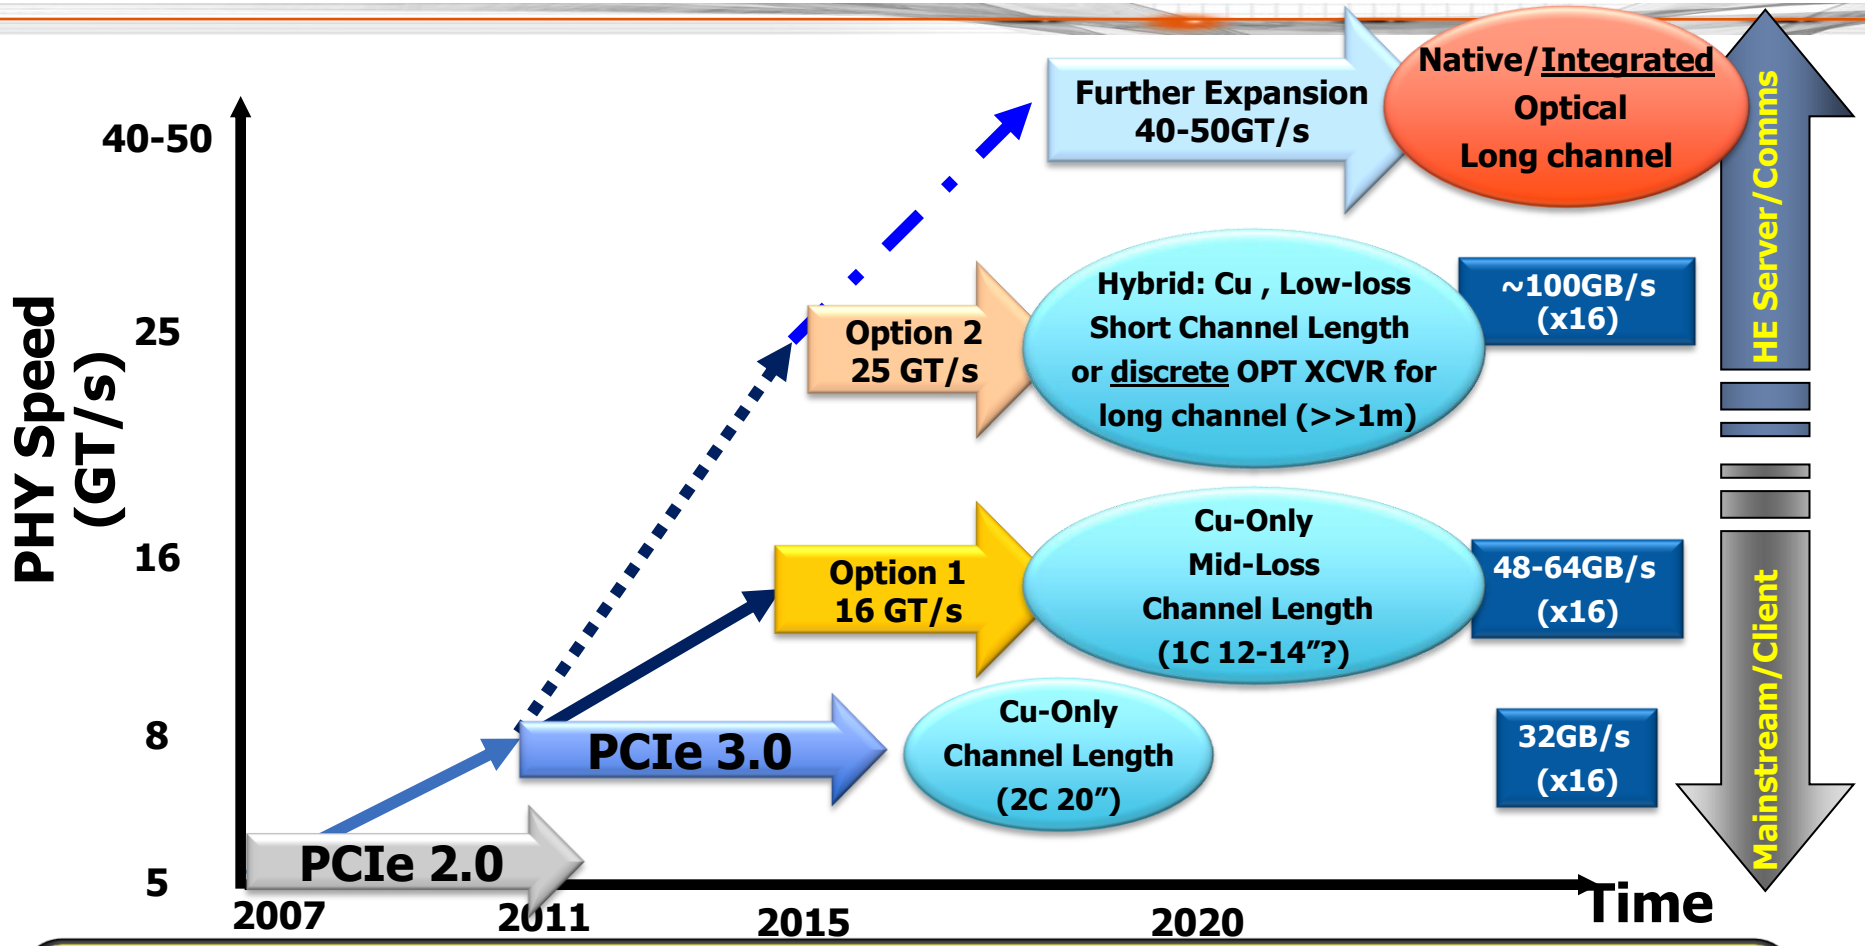

- *Option 1: Extending current usage models of PCIe*
- *Option 2: (1) Direct/short Cu interconnect to other IO/CPU and (2) Discrete optical XCVR for longer channel lengths*
- *Also looking at non-traditional interconnects*
- *These are from Intel Path-finding efforts*

Your Feedback Please…

OPENFABRICS ALLIANCE

- Data Rates
- Channels: disruption in form-factor? Riser/ Backplane?
- Backwards Compatibility
- Power Targets
- Power Management
- Protocol Enhancements

Further Comments to Debendra.Das.Sharma@intel.com